Beat Tracking with a Nonlinear Oscillator

Edward W. Large
Armstrong Laboratory AL/CFHP
2255 H Street, Bldg 248
Wright-Patterson AFB, OH 45433-7022
large@cis.ohio-state.edu

Abstract

The perception of beat and meter is a fundamental musical ability, yet machine simulation of this form of musical "common sense" remains a formidable AI problem. This paper outlines a nonlinear oscillator for beat perception in complex, metrically structured rhythms that has been described in detail elsewhere [Large, 1994; Large & Kolen, 1994]. A study is described in which pianists performed notated melodies and improvised variations on these same melodies. The performances are analyzed in terms of amount of rubato and rhythmic complexity, and the oscillator's ability to track beats in these melodies is assessed.

1 Introduction

The perception of beat and meter is a fundamental musical ability. This form of musical "common sense" enables apparently simple behaviors including tapping along with a tune, and very complex behaviors including the ability of skilled musicians to coordinate intricate motor activities with perceived musical rhythms and to notate musical performances. The perception of beat and meter by machine, however, remains a formidable problem. The difficulty of identifying temporal structure in complex signals such as music arises from several sources. One source of difficulty is the presence of *systematic timing deviations*. Performers use temporal deviation, or rubato, to communicate musical intentions, and such systematic deviations produce nonstationary input signals. Another source of difficulty is *rhythmic complexity*, which includes factors such as amount of syncopation and number of different duration values present in a rhythm. In short, the periodic components of rhythms that correspond to perceived beats are not really periodic, and even in ideally timed rhythms there are missing events and extraneous events.

Several approaches to the problem of beat tracking have been proposed [e.g. Allen & Dannenberg, 1989; Dannenberg, 1984; Dannenberg & Mont-Reynaud, 1987; Longuet-Higgins, 1987; Page, 1993; Rosenthal, 1992; Vercoe & Puckette, 1985]. For example, Dannenberg and Mont-Reynaud [1987] describe a history mechanism that uses a weighted average of previous perceived tempos to compute current perceived tempo. Allen and Dannenberg [1989] use a state description that includes phase and period, and real-time beam search to allow the beat-tracker to consider several possible states at once. Desain and Honing [1991] developed a connectionist quantizer to automatically "clean up" messy timing data in performance, and this approach has been extended to a create model of beat tracking and temporal expectancy.

This paper outlines a mechanism for beat perception in complex, metrically structured rhythms that has been described in detail elsewhere [Large, 1994; Large & Kolen, 1994]. The mechanism works on-line with local information. It possesses a memory for recent events, displays expectations for upcoming events, and can handle missing events at those times. The mechanism can also ignore events that should not affect its behavior. These properties are achieved by entraining a nonlinear oscillator to an incoming signal. Next, a study is described in which pianists performed notated melodies and improvised variations on these same melodies. The performances are analyzed in terms of amount of rubato and rhythmic complexity. The oscillator's ability to track beats in these melodies is assessed.

2 The Oscillator Model

Events (note onsets) in an input signal are represented as discrete impulses, $s(t)$. Figure 1 shows a series of impulses, corresponding to note events in an improvised melody, collected on a computer-monitored piano. $s(t) = 1$ when an event occurs, and $s(t) = 0$ at other times. The rhythmic input signal serves as a *driver*, and impulses in the signal perturb both the phase and the period of a *driven nonlinear oscillator*, causing changes to the oscillator's behavior. The oscillator adjusts its phase and period only at certain times during its cycle, isolating and tracking a periodicity in the incoming rhythm.

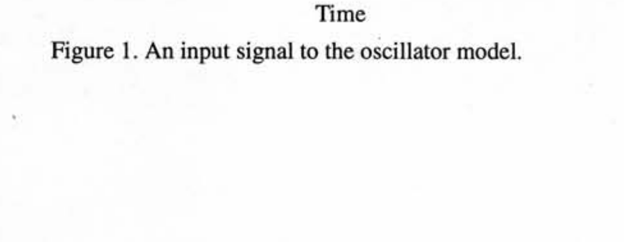

Figure 1. An input signal to the oscillator model.

2.1 Modeling Beats

The perception of beat is modeled by the generation an event at a particular phase of the oscillator's cycle. *Phase* refers to the specific point of the rhythm's cycle at a given time, t, defined at $-p/2 \le t < p/2$ as $\phi(t) = t/p$, where p is the period of the oscillation. According to this definition, phase varies between -0.5 and 0.5. The point $\phi(t) = 0$ is called t_x. t_x is the time at which the oscillator "expects" an event to occur. When an event occurs at time $t < t_x$, it is said to be early; when an event occurs at time $t > t_x$ is it said to be late. Thus, phase is:

$$\phi(t) = \frac{t - t_x}{p}, \qquad t_x - \frac{p}{2} \le t < t_x + \frac{p}{2}. \qquad \text{(Eqn 1)}$$

The events generated by the oscillator are called *output pulses*. An output pulse is like a music-theoretic beat except that is has a width, an extent in time. An output pulse is defined as:

$$o(t) = 1 + \tanh\gamma(\cos 2\pi\phi(t) - 1) \qquad \text{(Eqn 2)}$$

where γ is a parameter called *output gain* [Large & Kolen, 1994]. Figure 2 shows output pulses for two values of γ. Amplitude is non-zero for a relatively small portion of the cycle, defining a temporal receptive field for the oscillator. The parameter γ, the output gain, determines the size of the receptive field. The oscillator entrains to the signal by adjusting its phase and period only in response to signal impulses that occur within this receptive field; it ignores impulses that occur outside of this field. When γ is small, the receptive field is wide and the oscillator will tolerate more variability in the input signal. When γ is large, the region is narrow and the unit will tolerate less variability in the signal.

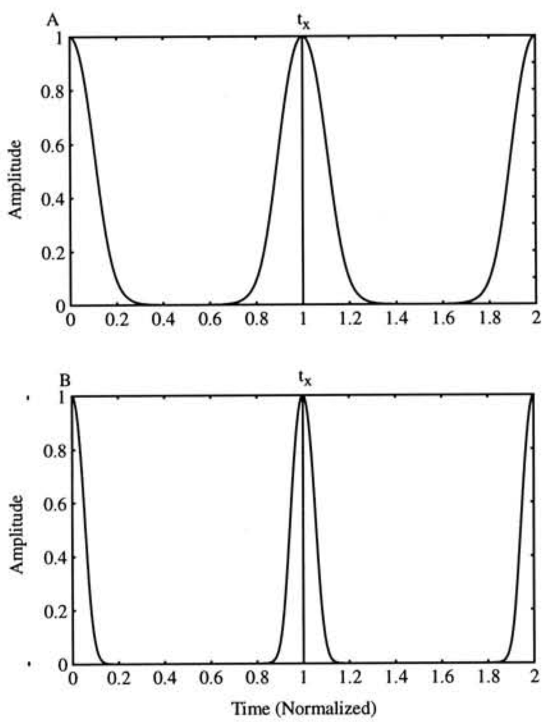

Figure 2. Output pulses (temporal receptive fields) for (A) $\gamma = 2$, (B) $\gamma = 8$.

2.2 Tracking Phase, Period, and Variability

To track an input signal, the oscillator adjusts its phase and period in response to input impulses. These behaviors are specified by delta rules that describe the nonlinear coupling between the input rhythm and the oscillator. These rules control adjustments to the phase and/or period in response to an impulse [Large, 1994; Large & Kolen, 1994]. Phase-tracking behavior can be implemented with the following rule:

$$\Delta t_x = \eta_1 s(t) \frac{p}{2\pi} \operatorname{sech}^2\gamma(\cos 2\pi\phi(t) - 1)\sin 2\pi\phi(t), \text{(Eqn 3)}$$

and period-tracking is implemented with a similar rule:

$$\Delta p = \eta_2 s(t) \frac{p}{2\pi} \operatorname{sech}^2\gamma(\cos 2\pi\phi(t) - 1)\sin 2\pi\phi(t). \text{ (Eqn 4)}$$

Here η_1 and η_2 are *coupling strength* parameters for phase and period tracking, respectively. These rules are related to the first time derivative of Equation 2. The relative extrema correspond to points of inflection on the output pulse curve [Large, 1994]. The presence of $s(t)$ in this formula ensures that adjustments will occur only when a signal impulse is present ($s(t) > 0$). An impulse that occurs within the oscillator's temporal receptive field, but before t_x causes a negative phase shift and a shortening of oscillator period. An impulse after t_x causes a positive phase shift and a lengthening of the period [Large, 1994].

The oscillator also adjusts the size of its temporal receptive field. This allows the unit to adapt to the amount of variability in the input signal. To do this, a third delta rule adjusts γ, through a control parameter, Ω, where

$$\Delta\Omega = -\eta_3 s(t) \operatorname{sech}^2\gamma(\cos 2\pi\phi(t) - 1) \cdot (\cos 2\pi\phi(t) + 2\gamma(o(t) - 1)\sin^2 2\pi\phi(t)). \text{ (Eqn 5)}$$

Here η_3 is the coupling strength for γ. This delta rule is related to the second time derivative of Equation 2. Zero-crossings correspond to the relative extrema of the phase- and period-tracking delta rules, giving this rule the power to adapt the oscillator's tracking behavior according to the variability of the input signal.

2.3 Decay, Confidence, and Compound Units

Because impulses may not occur in every cycle, γ decays toward 0 each cycle. Thus, if there is no event in the oscillator's current cycle, γ will decrease, widening the temporal receptive field.

Next, it is useful for the oscillator to adjust the amplitude of its output pulses, providing an internal measure of performance. Confidence, c, is used to measure the success of the oscillator in finding a periodicity in the input signal. There are a number of possible ways to do this. For the experiments that follow c is limited to a fixed range between $c_{min} = -1$ and $c_{max} = 1$, and increases monotonically with γ. Thus, as variability in the input signal shrinks, confidence grows.

Finally, in the examples that follow, each single unit is implemented using two oscillators operating in a tight, winner-take-all interaction. This arrangement is referred to as a *compound unit*. One oscillator, the *shadow* unit, is constrained to have the same period as the other, the *control* unit, but to remain exactly 180° out of phase. The unit with the greatest confidence at any given time is defined to be the control unit. The control unit controls phase and period adjustments and produces output pulses. However, both units (shadow and control) actively adjust γ, also controlling c. If at any time the shadow unit's confidence grows greater than that of the control unit, the two oscillators switch roles: the shadow becomes the control, and the control becomes the shadow. The observed behavior of a compound unit in this situation is a sudden 180° phase shift. This phase shift corresponds to a *gestalt* perceptual shift in the perception of the input signal – events that were perceived as out of phase with the beat are suddenly perceived as in phase. Figure 3 illustrates a situation in which a *gestalt* perceptual shift may occur, at $t = 3600ms$. Panel (D) in this figure graphs confidence of the control unit as a solid line, and confidence of the shadow unit as a dashed line. The gestalt perceptual shift happens when the two curves intersect. It also observable as a drop in cycle time, effective for a single cycle. Intrinsic period, p, is not affected. The usefulness of the arrangement will become apparent as specific examples are investigated.

3 Tracking Beats in Performances of Notated Melodies and Improvised Variations

3.1 The Test Data Set

To test the oscillator's ability to track beats in musical performances, a test set of sixty melodies was collected as follows [for detailed description of the data collection procedure, see Large, Palmer, & Pollack, 1995]. Two pianists performed melodies on a computer-monitored Yamaha Disklavier acoustic upright piano. Three children's melodies were chosen as performance material. For each melody, the pianists performed and recorded the melody, as presented in musical notation, five times. With the musical notation remaining in place, the pianists were then asked to play five improvisations. All performances were of a single-line melody only; pianists were instructed not to play harmonic accompaniment.

Next, skilled analysts transcribed the improvisations in standard musical notation and agreed upon the transcriptions. To assess the amount and distribution of rubato among the performances, a mean timing deviation was calculated for each performance, based either on the notation or the transcriptions. This measure of deviation was then averaged across the five performances of each melody or improvisation by each pianist. An analysis of variance (ANOVA) on mean rubato by performance type (notated melody vs. improvised variation), subject, and tune was conducted. There was a significant main effect of performance type ($F(1,4) = 33.46$, $p < 0.01$), indicating that, on average, more rubato was used in the improvisation of variations than in the performance of the melodies from notation. Mean rubato was 0.05 for notated melodies, and 0.10 for improvisations. There was also a significant interaction between tune and subject ($F(2, 8) = 13.89$, $p < 0.01$). Pianist 1 performed the melodies and improvisations for the first two tunes with little rubato, but for the third tune with high rubato. Pianist 2 performed tune three with little rubato, and performed tunes one and two with relatively high rubato.

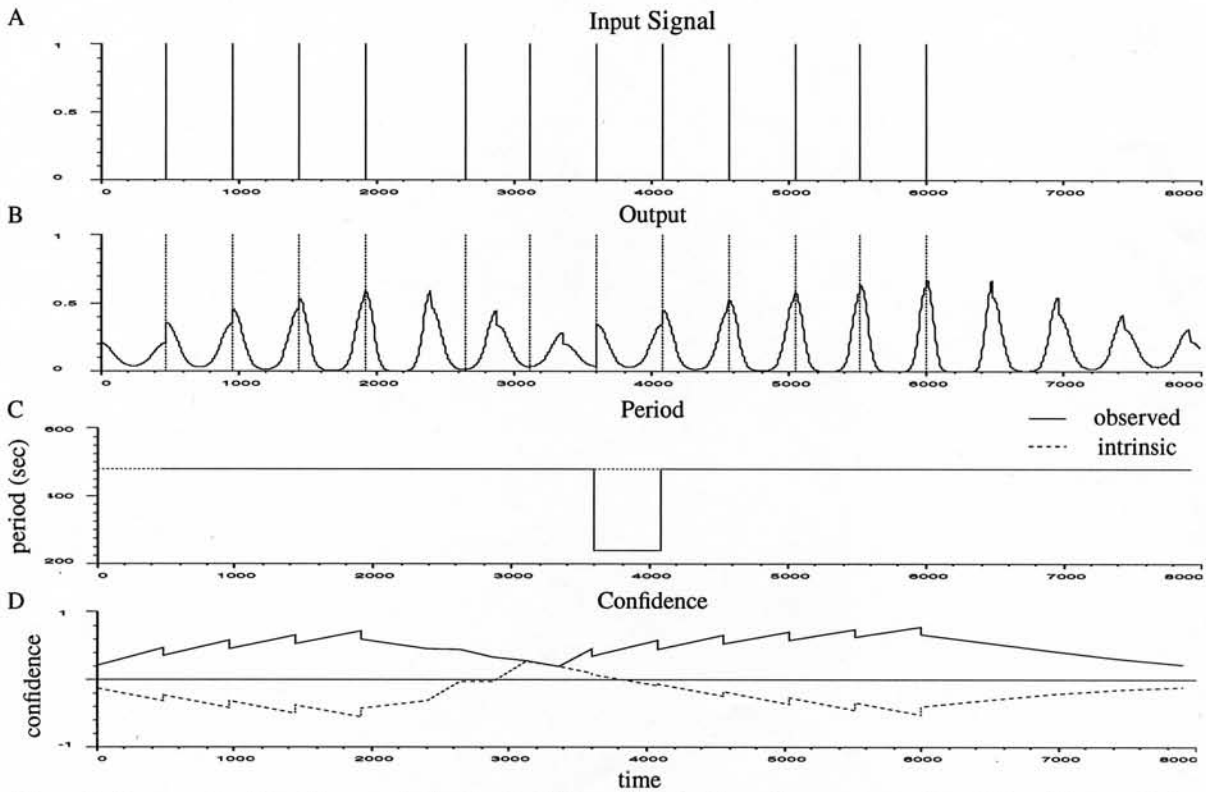

Figure 3. Compound unit tracking a periodic signal: (A) input signal, (B) oscillator output pulses, (C) oscillator period, (D) confidence.

3.2 Beat Tracking Performance

Next, the oscillator's ability to track beats in these melodies was assessed. The response of the oscillator was intended to model the perception of beats at some level in a metrical structure grid [Lerdahl & Jackendoff, 1983]. For each performance, the modal inter-onset interval (IOI) category was determined from the score or transcription, and chosen as the target metrical level. For each performance, the unit was initialized such that $\phi = 0$ at the initial onset, and p was set to the initial IOI of the target metrical level. Thus, the oscillator did not have to cope with finding initial phase or period.

A number of statistical measures of beat-tracking performance were collected. Only average absolute value of relative phase, $\overline{|\phi|}$, is reported here because this measure corresponded most closely to intuitive impressions of beat-tracking success, gleaned by listening to oscillator output. $\overline{|\phi|}$ measures performance as the average absolute phase of the oscillator on which events marking the beat occurred (as determined by the scores and transcription). $\overline{|\phi|} = 0$ means perfect performance, while $\overline{|\phi|} = 0.5$ is the poorest performance possible (meaning that the unit was 180° out of phase with the beat throughout the melody), and $\overline{|\phi|} < 0.1$ generally corresponded to a subjective impression of good performance. For each melody, oscillator performance was assessed by comparing $\overline{|\phi|}$ with the mean rubato score.

Performances of notated melodies and performances of improvised variations differed qualitatively in level of rhythmic complexity, and differed significantly in the magnitude of timing deviations, so results are discussed separately.

3.2.1 Performances of Notated Melodies

First, the oscillator was exposed to the thirty performances of notated melodies. Performances of notated melodies provided a controlled level of rhythmic complexity. Each melody contained three intended duration categories: sixteenth note, eighth note, and quarter note. Statistics were collected, and an analysis of variance (ANOVA) was conducted with factors tune, subject, and analysis type (mean rubato vs. average absolute phase). The ANOVA showed a main effect of analysis type ($F(1, 4) = 27.73$, $p < 0.01$), with mean rubato = 0.05, and average phase = 0.06. Thus, for these performances oscillator performed slightly worse than mean rubato would predict. This value of $\overline{|\phi|}$, however, indicates that on average the oscillator tracked the beat well. The ANOVA also showed a significant main effect of subject ($F(1, 4) = 11.19$, $p < 0.05$) indicating that some performers were more difficult to track than others. Therefore individual cases were examined in more detail.

Figure 4 gives an example of the oscillator's behavior as it tracks a performance of *Baa baa black sheep*. The first panel A provides a notated version of the melody (transcriptions of improvisations do not include grace notes or other ornaments) and a single row of dots from a metrical structure grid [Lerdahl & Jackendoff, 1983] marking the target events. Notes that are not marked by dots are to be ignored by the oscillator; dots that do not correspond to notes mark times when events are "missing" from the target event train. Panel B shows both input and output of the oscillator. The dashed lines show impulses in the input signal (marking event onset times). Because of the scale, full output pulses are not shown,

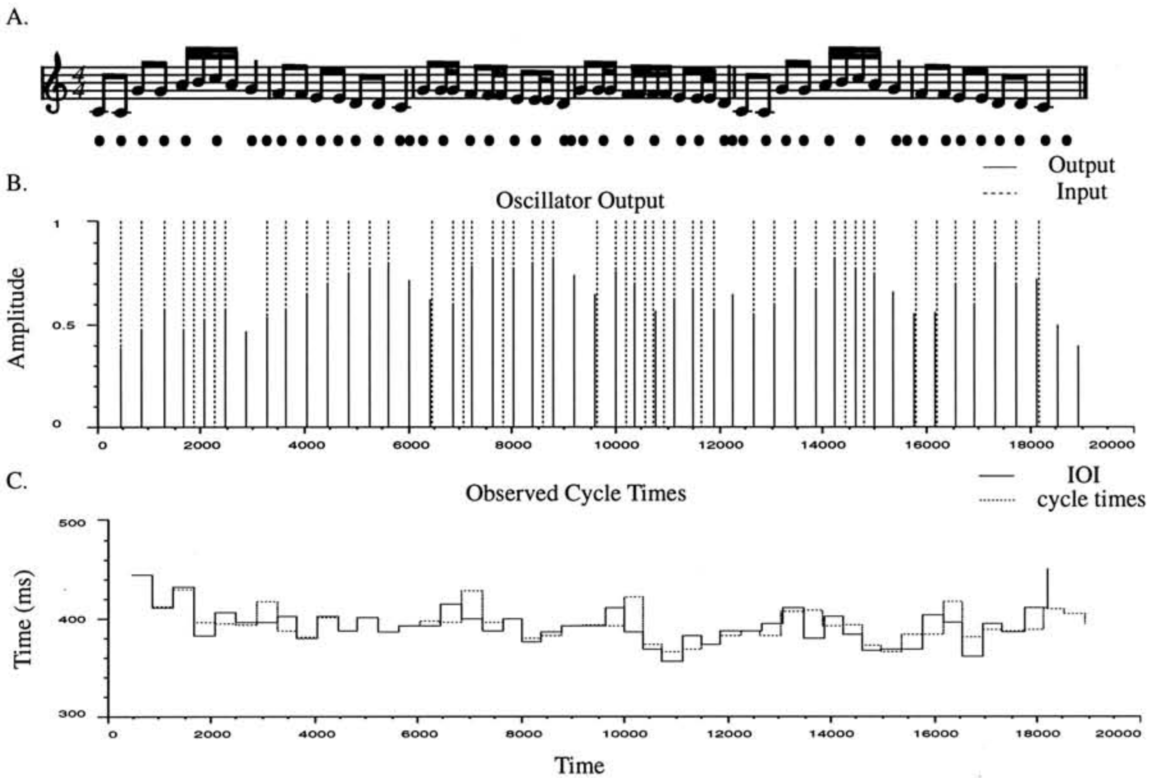

Figure 4. An oscillator tracking the rhythm of *Baa baa black sheep* (rubato = 0.05, $\overline{|\phi|} = 0.08$).

rather discrete output pulses (shown as solid lines) are displayed at $t = t_x$. Discrete output pulses are used for producing auditory output via MIDI, so this panel describes what the oscillator's output sounds like. These two lines overlap when a target event is performed at precisely the time predicted by the oscillator, that is, at phase zero, $\phi(t) = 0$, of the driven oscillator. Amplitude of the oscillator output is controlled by confidence, c, providing a way for the oscillator to measure its own performance. High amplitude of the discrete pulse corresponds to a small temporal receptive field, low amplitude corresponds to a wide receptive field.

Panel C shows a tempo curve for the performance as a solid line. This curve was derived by extracting the target events from the performance and graphing IOI's for these events. This curve gives the IOI's to which the oscillator should respond. Panel C shows actual observed cycle times of the oscillator using a dotted line. Observed cycle time takes into account not only the intrinsic period, p, of the oscillator, but also phase t_x as it is adjusted in each cycle. Thus, this curve represents the combined effect of the two delta rules given by Equation 3 and Equation 4. Beginning at the initial tempo, the unit effectively calculates a local tempo, and follows performance tempo as the performer speeds up and slows down, based on eighth note onset times only.

This is an example of the oscillator behaving quite well. Case by case examination revealed that 29 out of the 30 cases matched this profile, the oscillator tracked its target without difficulty ($\overline{|\phi|} < 0.10$). In 1 case, however, average phase was high (rubato = 0.07, $\overline{|\phi|} = 0.20$). The largest value that $\overline{|\phi|}$ can assume is 0.5, when a target event onset occurs at anti-phase with the oscillator's pulses. Thus, this value of average phase indicates that the unit had some difficulty tracking the target events.

Figure 5 shows what the difficulty was. At the end of the first half of the melody, the oscillator has done a good job of estimating period (panel C), but the first note of the second phrase comes in at nearly anti-phase from the oscillator pulses. For the next three cycles, the oscillator preserves this anti-phase relationship with the target event train; this is the source of the high value for average phase. On the third cycle, the oscillator performs a *gestalt* perceptual shift (a sudden 180° change in phase, described above). The perceptual shift can be seen as a dramatic drop in observed cycle time in panel C. After the shift, the oscillator successfully tracks the remainder of the performance.

Timing deviations in which the performer enters 180° out of phase from the output pulses present a common difficulty for the oscillator. The difficulty arises because target events occur at times when the oscillator has decided to ignore event onsets. The unit is able to recover because it can perform a perceptual shift. The value of γ decays, accompanied by a drop in confidence, allowing the shadow unit to take over. This presents a difficult situation for the oscillator, but one from which it can quickly and gracefully recover.

3.2.2 Improvisation of Variations

Next, the oscillator's performance in tracking the thirty improvised variations was examined. The improvisations provided a more difficult situation than the performances of notated melodies for two reasons. The rhythms of the improvisations were more complex than the rhythms of the

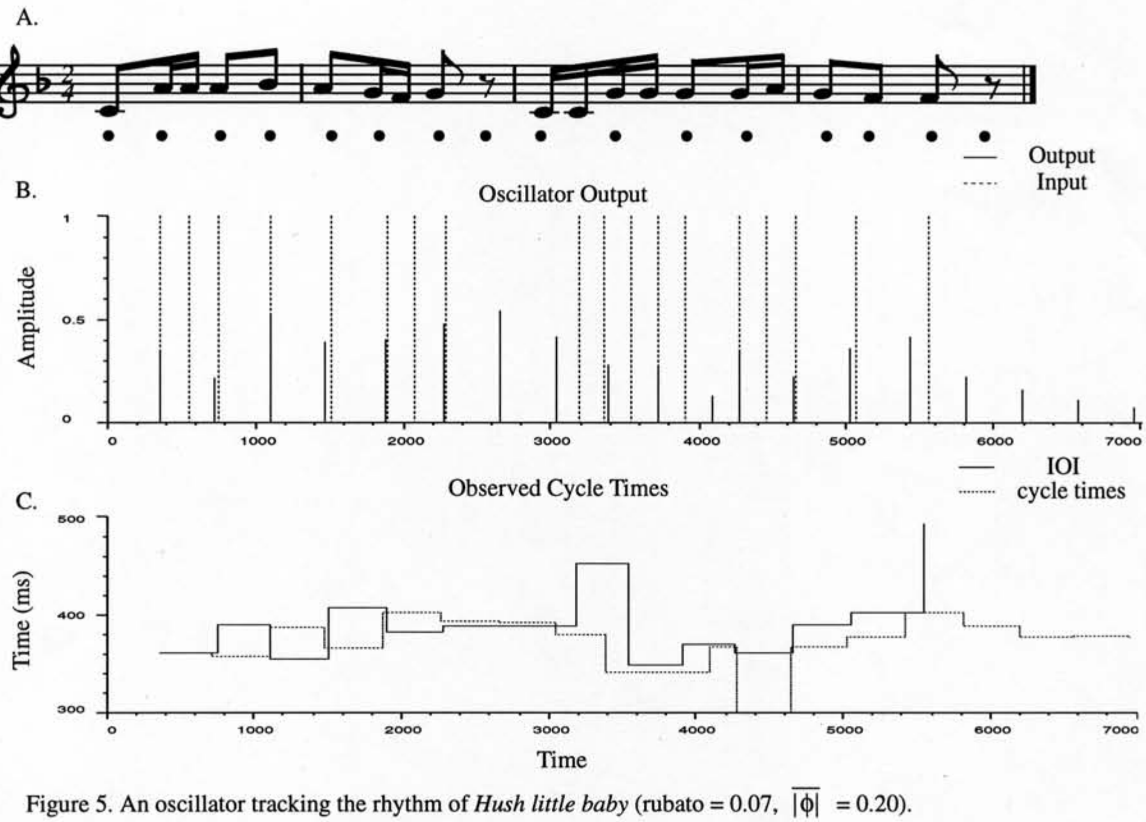

Figure 5. An oscillator tracking the rhythm of *Hush little baby* (rubato = 0.07, $\overline{|\phi|} = 0.20$).

melodies, making use of syncopation, and containing up to seven different levels of intended durations according to the transcriptions. Also the improvisations showed significantly greater timing deviations than did the performed melodies. The oscillator was exposed to the melodies, statistics were collected, and an analysis of variance (ANOVA) was conducted with factors tune, subject, and analysis type (mean rubato vs. average absolute phase). The ANOVA showed no main effect of analysis type ($F(1, 4) = 0.005$, $p = 0.947$), with mean rubato = 0.10, and average phase = 0.10. This result shows that for these performances, beat tracking performance is on par with mean rubato. The ANOVA also indicated a significant interaction of tune and subject ($F(2, 8) = 4.0$, $p < 0.05$), indicating that the oscillator had more trouble with some performances than with others.

Case by case examination reveled that in 20 out of the 30 cases, the oscillator tracked its target well ($\overline{|\phi|} < 0.10$). In 10 cases the unit had some difficulty in tracking beats. The 10 difficult cases were examined in detail. Two of the most difficult are discussed here.

Pianist 1's improvisations on *Mary had a little lamb* were performed in a freely timed blues style. The first improvisation had the highest rubato score (rubato = 0.25), highest average phase ($\overline{|\phi|} = 0.17$). The oscillator's behavior in this case was representative of its performance on this group of melodies, so it was chosen for further study. The time series corresponding to the performance of the oscillator are shown in Figure 6. The tempo curve indicates the presence of large timing deviations at several points in the melody. Points of particular interest are around t = 3000*ms*, t= 8000*ms*, t= 16000*ms*, and t = 23000*ms*. At these points, γ drops (correspondingly, confidence drops) allowing the oscillator to continue to track the target in spite of the large deviations. In three of these cases the compound unit responds to difficult timings with perceptual shifts.

In spite of these difficulties, however, the figure shows that the oscillator did a respectable job of tracking its target event train in this rhythm. Beats are output at approximately the correct times throughout the piece – the oscillator is not lured away by the many distractor events in this rhythmically complex performance. Another way to see this is to note that the value of average absolute phase ($\overline{|\phi|} = 0.17$) is lower than mean rubato (0.25). Additionally, oscillator confidence is high for large sections of the piece; by the oscillator's internal measure its performance is good.

9 of the 10 problem cases examined fit the profile of this case. The oscillator had difficulty in certain sections of the performances, but always recovered gracefully. In two of these cases, the difficulties were caused by actual timing errors (what appeared to be "stumbles" and/or "think time") on the part of the performers [Large, 1994]. In all cases, oscillator confidence was high throughout most of the performances.

In 1 of the 10 difficult cases examined, however, the result was poor. Pianist 2's improvisations on *Hush little baby* were the most varied of all the improvisations studied and made heavy use of rubato. The improvisation that proved the most difficult for the model to handle was the third variation. This improvisation made moderately heavy use of rubato and had the highest average phase, (rubato = 0.16, $\overline{|\phi|} = 0.30$). These

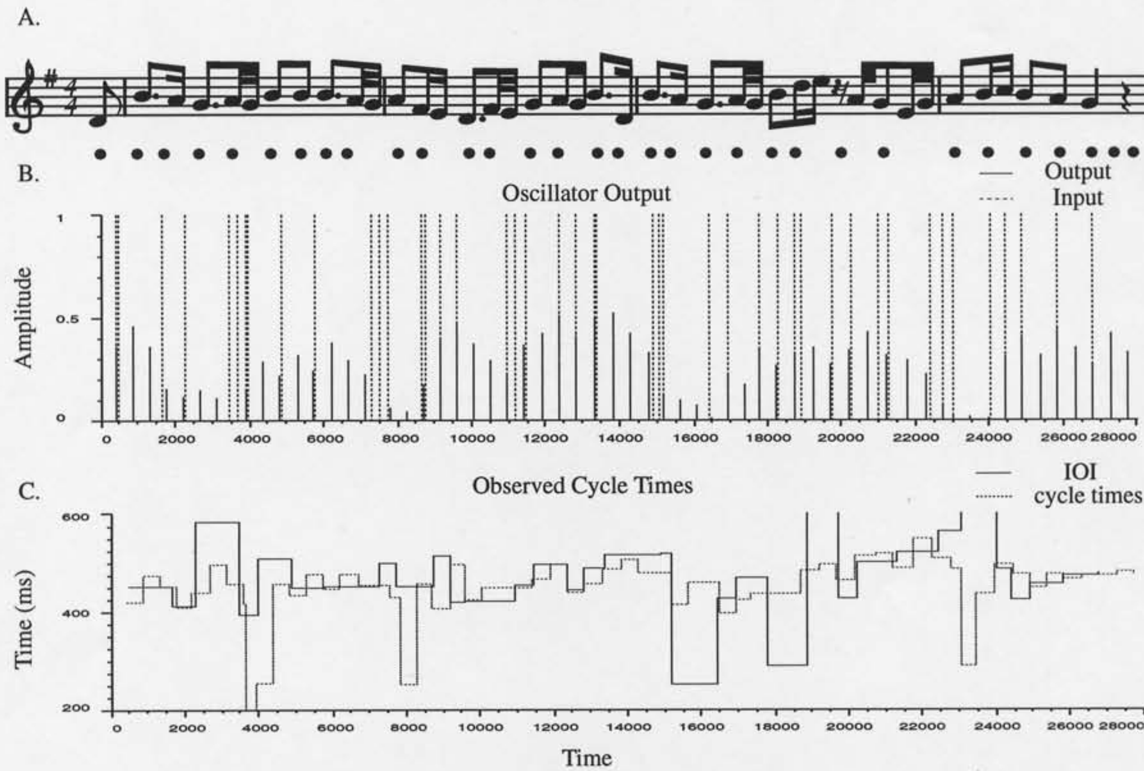

Figure 6. Oscillator tracking an improvisation on *Mary had a little lamb* (grace notes are not transcribed) (rubato = 0.25, $|\phi| = 0.17$).

numbers suggest extreme difficulty in tracking. Figure 7 shows the actual time series corresponding to the performance of the oscillator. Throughout this improvisation the performer makes use of a sort of "jagged" rubato. The tempo curve, shown in panel C, reveals timing deviations that strictly alternate: slower, faster, slower, faster. Panel C also shows the effect of this pattern on observed cycle times. Cycle times are always one step behind the performed durations because changes to the oscillator's phase and period in the current cycle effect oscillator cycle time for the following cycle.

Because timing deviations zigzagged in this fashion, cycle time decreased when performed duration increased, and vice-versa. This rubato pattern occurred in other performances as well, however in this case the amount of rubato was large enough to pose a serious difficulty for the oscillator. The event at time $t = 1800ms$ is very early (almost anti-phase) the confidence of the unit decreases. By the downbeat of the second full measure, a perceptual shift takes place. This situation repeats itself until nearly the end of the performance. Toward the end the performer regulates the timing, and the unit finally begins to pick up confidence.

The end result was that the oscillator tracked this performance in an odd way. The unit responds to the correct events, and outputs beats at more-or-less the correct locations throughout the piece. However, the oscillator's confidence is consistently low because there is so much temporal deviation. Thus large, jagged rubato patterns represent a limiting case for the single oscillator model.

4 Discussion

Melodies are perhaps the most difficult cases for beat tracking systems, because they provide fewer reliable cues than accompanied melodies. Bass lines and harmonic accompaniment tend to be more rhythmically consistent, providing additional information for the system. In this study, performances of notated melodies provided a controlled level of rhythmic complexity, while improvised variations provided syncopation and a great variety of duration categories. Both types of performance contained timing deviations, making the task of tracking a single periodicity a challenging one. Yet, in 49 out of 60 cases, the oscillator tracked performed rhythms well by an objective measure ($\overline{|\phi|}$).

In 11 cases, difficulties were encountered ($\overline{|\phi|} > 0.10$). These difficulties were caused by large temporal deviations, stemming from three sources: heavy use of rubato including 'phase-shifts', actual timing errors on the part of performers, and jagged rubato curves resulting from alternating shortened and lengthened durations [Large, 1994]. In 10 of the 11 cases, however, the oscillator tracked nicely for large sections of the melodies, having trouble in some areas but recovering gracefully.

In the most difficult case, an improvisation that coupled heavy rubato with a jagged tempo profile, the oscillator performed poorly $\overline{|\phi|} = 0.30$. Although it did not lose the beat altogether, it was not able to adequately follow the tempo changes, and its internal measure of performance was consistently low. This case illustrates the limits of a single oscillator model. All may not be lost, however. Note that

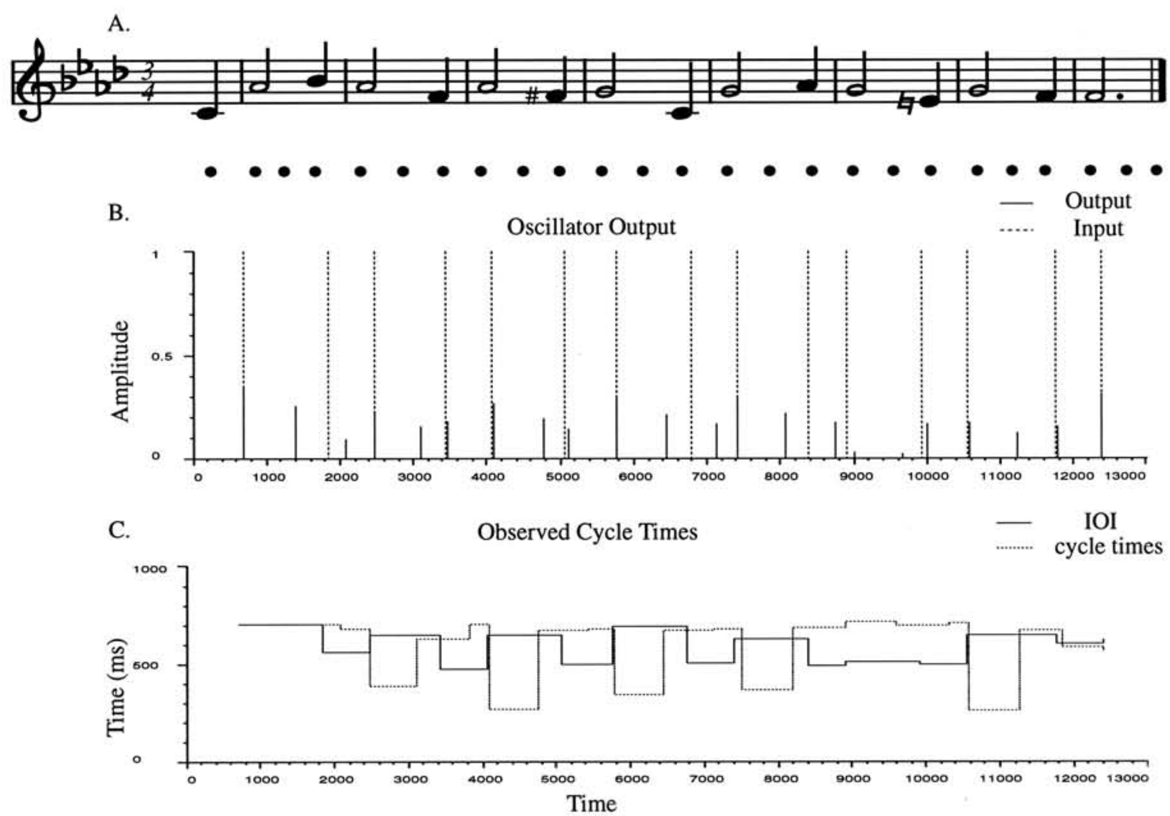

Figure 7. Oscillator tracking an improvisation on *Hush little baby* (rubato = 0.16, $\overline{|\phi|} = 0.30$).

tempo changes of approximately the same magnitude strictly alternate (Figure 7, Panel C). This means that had the oscillator been attempting to track at the dotted half-note level, it would have found almost no rubato at all. Cases such as this would tend to argue for a multiple oscillator model. In a multiple oscillator model different oscillators track different levels of beats, as in a metrical structure grid [Large & Kolen, 1994; Lerdahl & Jackendoff, 1983]. Internal interactions phase-lock the oscillators, and the perception of metrical structure can emerge from the perception of beat on multiple time scales. Building networks of interacting oscillators to model the perception of metrical structure is a focus of current research.

Overall, the single oscillator model did a remarkable job of tracking the beat in complex rhythms with no information other than event onset times. These analyses suggest that nonlinear oscillators, driven with complex, nonstationary rhythms that arise from musical performance, can adequately model the perception of musical beat, and may ultimately be used to model the perception of musical meter.

5 References

[Allen and Dannenberg, 1989] P. E. Allen and R. B. Dannenberg. *Tracking musical beats in real time*. In Proceedings of the 1990 International Computer Music Conference (pp. 140-143). Computer Music Association.

[Allen and Dannenberg, 1989] R. B. Dannenberg. *An on-line algorithm for real-time accompaniment*. In Proceedings of the 1984 International Computer Music Conference. Computer Music Association.

[Dannenberg and Mont-Reynaud, 1987] R. B. Dannenberg and B. Mont-Reynaud (1987). *Following an improvisation in real time*. In Proceedings of the 1987 International Computer Music Conference. Computer Music Association.

[Desain and Honing, 1991] P. Desain and H. Honing. The quantization of musical time: A connectionist approach. In P. M. Todd and D. G. Loy (Eds) *Music and Connectionism* (pp. 150-160). Cambridge: MIT Press.

[Large, 1994] E. W. Large. *Dynamic representation of musical structure*. Unpublished Ph.D. Dissertation. The Ohio State University, Columbus, OH.

[Large and Kolen, 1994] E. W. Large and J. F. Kolen. Resonance and the perception of musical meter. *Connection Science, 6 (1)*, 177 - 208.

[Large, Palmer, and Pollack, 1995] E. W. Large, C. Palmer, and J. B. Pollack. Reduced memory representations for music. *Cognitive Science, 19*, 53-96.

[Lerdahl and Jackendoff, 1983] F. Lerdahl and R. Jackendoff. *A generative theory of tonal music*. Cambridge: MIT Press.

[Longuet-Higgins, 1987] H. C. Longuet-Higgins. *Mental Processes*. Cambridge: MIT Press.

[Page, 1993] M. P. A. Page (1993). Modelling aspects of music perception using self-organizing neural networks. Unpublished doctoral dissertation, University of Cardiff.

[Rosenthal, 1992] D. Rosenthal. Emulation of human rhythm perception. *Computer Music Journal*, 16, 64-76.

[Vercoe and Puckette, 1985] B. Vercoe and M. Puckette. *Synthetic rehearsal: Training the synthetic performer.* In Proceedings of the 1985 International Computer Music Conference (pp. 275-278). Computer Music Association.